arXiv:1611.00252v1 [cs.LG] 30 Oct 2016

Improving a Credit Scoring Model by Incorporating Bank Statement Derived Features

Rory P. Bunker, Wenjun Zhang, M. Asif Naeem

Auckland University of Technology

Auckland, New Zealand

Abstract

In this paper, we investigate the extent to which features derived from bank statements provided by loan applicants, and which are not declared on an application form, can enhance a credit scoring model for a New Zealand lending company. Exploring the potential of such information to improve credit scoring models in this manner has not been studied previously. We construct a baseline model based solely on the existing scoring features obtained from the loan application form, and a second baseline model based solely on the new bank statement-derived features. A combined feature model is then created by augmenting the application form features with the new bank statement derived features. Our experimental results using ROC analysis show that a combined feature model performs better than both of the two baseline models, and show that a number of the bank statement-derived features have value in improving the credit scoring model. The target data set used for modelling was highly imbalanced, and Naive Bayes was found to be the best performing model, and outperformed a number of other classifiers commonly used in credit scoring, suggesting its potential for future use on highly imbalanced data sets.

Keywords: Credit Risk, ROC Curve, Imbalanced Data, Machine Learning

1. Introduction

Credit scoring models are widely used by banks and other financial institutions, in order to assess the risk of default of applicants for loans. Credit scoring can be thought of as a a classification problem. Credit scoring models take a vector of attributes for a loan applicant, and given these attributes,

attempt to discriminate between goods and bads; that is, to discriminate between those that are not likely to default or be in arrears with their payments, and those that are.

A lending company in New Zealand scores loan applicants based on 11 attributes obtained from an application form, which is filled out online when applying for a loan. The lending company had recently begun to incorporate data from an application called Credit Sense (https://creditsense.com.au) into their data warehouse, which automatically extracts line-by-line bank statement data from loan applicants bank statements for a 90-day period of spending. The bank statement data is currently used for purposes such as income verification. The company was interested in exploring the potential of this data to be used for credit scoring.

The company was firstly interested in seeing whether a sufficiently predictive scoring model could be created using only the bank statement-derived features, in which case the questions in the online application form would then be unnecessary to ask. If there are less questions to ask, loan applicants would be able to complete the application process faster, and the lending company would potentially be able to write more loans. If this model was not found to be sufficiently predictive, then the lending company was interested in investigating the extent to which the bank statement derived features could be of use in developing an improved scoring model by supplementing existing scoring features.A number of machine learning and statistical models have been used for experimentation in this paper.

The novelty of our contribution is in the investigation of deriving features from bank statement data provided by loan applicants, which are not declared on the loan application form, to be used in credit scoring models. To the best of our knowledge, has not been studied previously. For this investigation, firstly we construct **two baseline models**. The **first**, based solely on the existing scoring features, and the **second**, based solely on the new bank statement derived features. Secondly, we create a **combined feature model** by combining the existing scoring features with the new bank statement-derived features.

Our experimental results using ROC analysis show firstly that the application form feature baseline model outperforms the bank statement baseline model, and in turn, a combined feature model performs better than both of the two baseline models. This indicates that the bank statement-derived features have value in supplementing, but not in replacing, the application form features. A Naive Bayes model consisting of the 11 existing application

form features, along with 5 new bank statement derived features, was found to be the best performing model.

This paper is organized as follows. Some background on credit scoring will be provided in section 2. Section 3 will outline the research goals for this paper. Section 4 will then provide a discussion of the methodology in particular, the knowledge discovery in databases framework that will be used to guide the data mining and knowledge discovery process. Section 5 will provide a brief outline of the candidate models that will be used in the experimentation. Section 6 will review metrics for the evaluation of classification models for imbalanced data sets, and section 7 will discuss the characteristics of the target data set used. Section 8 will summarize the experiments and their results with the different classifiers and feature sets. Section 9 will evaluate the final model, and section 10 will conclude and discuss some avenues for further work.

2. Credit Scoring

Thomas, Edelman & Crook [1] define credit scoring as "the set of decision models and their underlying techniques that aid lenders in the granting of consumer credit. These techniques decide who will get credit, how much credit they should get, and what operational strategies will enhance the profitability of the borrowers to the lenders". The present study is focused on models deciding who will get credit (application scoring), rather than on modelling the behaviour of existing loans (behavioural scoring).

There has been a significant amount of academic research in the area of credit scoring. There exists a number of review papers in the literature on credit scoring including that by Rosenberg and Gleit [2], Hand and Henley [3] and Thomas [4]. More recent expositions include that by Abdou and Pointon [5]. Yu et al [6] also provides a useful summary in chapter 1, althought the book primarily focused specifically on the use of Support Vector Machines for credit scoring.

There are a number of important benefits from making use of credit scoring models. While initially credit analysts in the industry used methods based on their own judgment, sophisticated statistical methods found their way into use and have generally been found to be more subjective, as well as enabling the large scale automation of the loan acceptance process, and writing more loans than they otherwise would have been able to, at lower cost. Economic pressures from an increase in the demand for credit, as well

as greater competition, and the development of new technology, are all factors that have led to the development of statistical models to aid the lending decision [3]. Not only are scoring models the only way of handling the large number of applicants nowadays, but also they seemingly produce more accurate classifications than the subjective judgemental assessments by human experts [2].

In assessing the credit risk of a loan applicant, a lending company typically obtains applicant characteristics via an application form, as well as credit bureau checks, which are institutions that check the past credit history of applicants i.e. whether they have paid off previous loans that they have taken out in the past, whether they have had any defaults, or whether they have been declared bankrupt and so forth. In some cases, an initial application form screens good and bad risk applicants, and only those who proceed past this stage are checked by credit bureau agencies therefore avoiding unnecessary costs and providing quick credit application decision time to most applicants [3]. In this case, the lending company sequentially assesses a loan application.

As has been mentioned, credit scoring can be viewed as a classification problem, where we attempt to discriminate between good and bad risks, given their characteristics. A useful mathematical summary of the credit scoring evaluation problem is outlined in [6], and is summarized below:

Let

$$\mathbf{x} = (x_1, x_2, , x_m)^T \tag{1}$$

be a vector of m random variables that describe the information from a customers application form and credit reference bureau. The actual value of the attributes for a specific loan applicant k are

$$\mathbf{x}_k = (x_{1k}, x_{2k}, , x_{mk})^T \tag{2}$$

All samples are then given by

$$S = (x_k, y_k), k = 1, 2, , N \tag{3}$$

where N is the number of loan applicant samples in the data set, x_k is the attribute vector of the kth applicant, and y_k is the observed result of whether or not there was timely repayment. For example, we could define $y_k = 1$ if the customer repaid the loan i.e. was a good, and $y_k = 0$ if applicant k did not repay the loan (or was in arrears for some time).

In cases where we have both the attribute vector and the observed loan outcome, the data is labeled and thus we are able to make use of supervised learning models. Therefore, it is natural to make use of machine learning and statistical models that are used for classification purposes.

In order to make an accurate judgment of the likelihood of y_k, given the applicants attribute vector $\mathbf{x}_k$, there have been a number of models that have been applied to the problem, drawing from different disciplines including machine learning, statistics and operations research. By analyzing the techniques used in a large number of research papers, [6] grouped these techniques broadly to fit within 5 categories (table 1).

Technique Category	**Technique**
Statistical Models	Linear Discriminant Analysis
	Logistic Regression
	Probit Regression
	K-Nearest Neighbors
	Decision Trees
Mathematical Programming	Linear Programming
	Quadratic Programming
	Integer Programming
Artificial Intelligence	Artificial Neural Networks (ANN)
	Support Vector Machines (SVM)
	Genetic Algorithm
	Genetic Programming
	Rough Set
Hybrid Approaches	ANN and Fuzzy Systems
	Rough Set and ANN
	Fuzzy System and SVM
Ensemble Approaches	ANN Ensemble
	SVM Ensemble
	Hybrid Ensemble

Table 1: Techniques for Credit Scoring [6]

While traditionally statistical and linear programming techniques were most commonly applied in credit scoring, sophisticated methods from artificial intelligence and machine learning have also become commonly applied in the area of credit scoring. Due to time constraints, a selection of the models

in table 1 will be used in the present study.

3. Research Goals

The main goals of the project are as follows:

1. Extract relevant features from the bank statement data that can be potentially used for modelling credit risk.
2. Construct baseline models:
    - based solely on the data set consisting solely on the existing scoring features, obtained from the information declared on the loan application form.
    - based solely on the data set consisting solely on the bank statement derived features, obtained from Credit Sense data provided by loan applicants.
3. By comparing the performance of the combined feature model (application form + bank statement-derived) with the two baseline models, investigate whether the bank statement derived features are sufficiently predictive by themselves for a credit scoring model, or the extent to which the bank statement-derived features can supplement the existing scoring features to enhance a credit scoring model.

4. Methodology

For this project, a structured experimental approach to the data mining process was sought. The knowledge discovery in databases process (KDP) outlines a formal experimental approach to extracting and creating knowledge from data. The KDP is defined as the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data [7]. The KDP involves having some input data, and in the end, coming out with some new knowledge that is valuable and useful to the organization.

Models of the KDP were initially developed in academia, and models from industry followed [9]. Two models that are considered to be leading ones are that developed by Fayyad et al in [7], which was developed from academic research, and the second is the Cross Industry Standard Process for Data Mining (CRISP-DM) model [23], which was developed out of a project led by a number of organizations from industry.

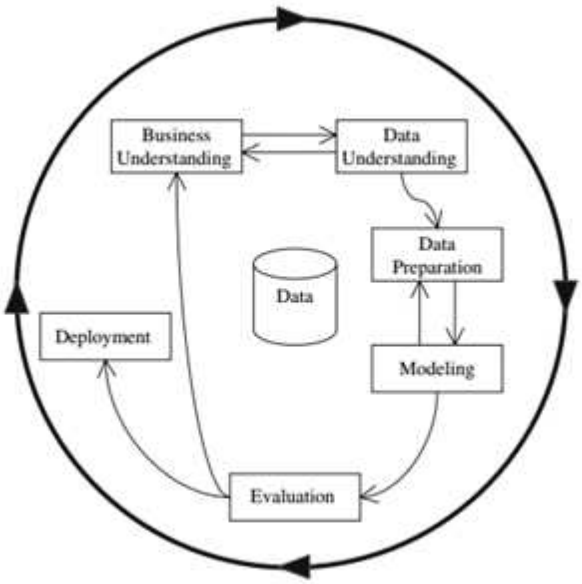

Figure 1: The CRISP-DM model of the KDP. Source: Krzysztof et al, 2007

We apply the CRISP-DM methodology in the present study, due to it being an industry-specific framework, and this research being performed in conjunction with industry. This framework consists of 6 steps:

In the first two steps, the data miner needs to gain understanding of the business or application domain. In the current project, this involves becoming familiar with the business and its data. This is a natural process of spending time in the company, talking with staff, as well as preliminary investigation and querying of the lending companys data warehouse.

Data preparation or preprocessing involves the (often time-consuming) process of creating a final data set that can be used as an input to the machine learning or statistical model. The data preprocessing indeed proved to be a time consuming aspect of the present study (extracting new features from the bank statement data, and augmenting these features with the existing scoring features to create a final target data set for modelling). Nonetheless, it is a very important step as it can considerably affect the success of the data mining experiments and helps to ensure that one does not obtain misleading results [8].

Algorithm selection in the present study involves investigating previous

literature in the application domain, and selecting and experimenting with algorithms that have been successful in previous studies. The experimentation with the different candidate models constitutes the data mining and modelling step.

The final steps are to evaluate the model, its potential use to the company, and to possibly deploy it in some way. This involves putting the results and/or knowledge obtained from the model to use e.g. through deploying the model into existing company IT systems, or simply in the production of a report outlining the analysis.

5. Candidate Models

An important step in the knowledge discovery process is selecting candidate models. The candidate models that were used for experimentation in this study are as follows:

- Logistic Regression
- Naive Bayes
- Support Vector Machine with a linear kernel
- Nearest Neighbor
- J48 Decision Tree
- Random Forest

Linear Discriminant analysis (LDA) was not considered in this study due to its difficulty in dealing with nominal features. We also decided not to consider mathematical programming approaches, as to consider these would require a completely different experimental set-up in terms of defining objective functions and constraints. Naive Bayes, interestingly, was not noted by [6] as being a common technique in credit scoring; however for comprehensiveness, it was included in this study.

For simplicity, all schemes were run in WEKA with their default parameter settings.

We now briefly outline some of the theoretical background of each of the candidate models that will be used in the experimentation.

5.1. Logistic Regression

Logistic Regression is a general linear model (GLM) that models a binary outcome (0/1, good/bad etc.) on a certain number of predictors. Logistic regression is the most commonly applied and a strong performing method for credit scoring, and has a strong theoretical basis in that it directly gives us an additive log odds score, which is a weighted linear sum of the attribute values [33].

A logistic transformation yields a dependent variable that is between 0 and 1, and this provides an intuitive interpretation, being the probability of default given the attribute vector for that particular loan applicant. The logistic function is given by

$$p(x) = \frac{e^x}{1+e^x} = \frac{1}{1+e^{-x}} \tag{4}$$

In logistic regression, we substitute in the standard linear regression model into the logistic function (4). That is,

$$y = \beta.\mathbf{x} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_k x_k \tag{5}$$

$$p(\mathbf{x}) = \frac{1}{1+e^{-\beta.\mathbf{x}}} \tag{6}$$

where y is the class value (good, bad), and $\mathbf{x}$ is a vector of applicant characteristics.

Rearranging (6) and taking the natural logarithm of both sides yields

$$\frac{p(\mathbf{x})}{1-p(\mathbf{x})} = e^{\beta_0+\beta_1 x_1+\beta_2 x_2+...+\beta_k x_k} \tag{7}$$

$$ln\frac{p(\mathbf{x})}{1-p(\mathbf{x})} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + ... + \beta_k x_k \tag{8}$$

If p is defined as taking the value 0 if the loan applicant is a bad risk, and 1 if the loan applicant is a good risk, then the term p($\mathbf{x}$) is the conditional probability of being a good risk, given the loan applicants attribute vector $\mathbf{x}$.

Estimation of the coefficients of the model in (8) can be achieved using maximum likelihood estimation (MLE), where we try to find estimates for the beta coefficients such that the predicted probability of default for each individual corresponds as closely as possible to the individuals observed status [14]. This MLE procedure is built into most modern statistical computing packages.

5.2. Naive Bayes

The Naive Bayes classifier is a probabilistic classifier that uses the well-known Bayes Theorem from probability theory. We follow the notation of [12], and state Bayes Theorem as

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)} \tag{9}$$

where H denotes the hypothesis (the class of the particular instance), and E represents the evidence (the attribute values). We would like to estimate the probability of the class for a particular instance, given the attribute vector. Given an applicant attribute vector $\mathbf{x}$ and a new instance to classify, if P($Good|x$) is larger than $P(Bad|x)$ then we classify the new instance as good, but otherwise we classify the new instance as bad. The key assumption of Naive Bayes is that the attributes are conditionally independent of each other, given the class. When we have conditional independence of attributes given the class value, we can write (9) as

$$P(H|E) = \frac{P(E_1|H)P(E_2|H)...P(E_n|H)P(H)}{P(E)} \tag{10}$$

Where P(H) is the prior probability of getting class H, regardless of the attributes, and the $P(E_i|H)$ denote the probability of obtaining that attribute value E_i, given the class value. Although we do not have the denominator in Bayes Theorem P(E), one can normalize the likelihoods so that the $P(H|E)$ probabilities add up to one for the different class values.

Although the conditional independence assumption rarely holds in practice, the Naive Bayes classifier has nonetheless often provides impressive results [12]. Its probabilistic foundation is also relatively interpretatble and intuitive.

5.3. Support Vector Machines

Support vector machines (SVM) perform classification by attempting to divide the points that belong to the two classes by constructing the largest possible 'wedge' between the two classes. The SVM was first proposed by [13].

The SVM produces nonlinear boundaries by constructing a linear boundary in a large, transformed version of the feature space (by using a kernel

function), and the method does not require that the classes be linearly separable [14]. A soft margin allows some observations to be on the incorrect side of the margin or hyperplane. A parameter known as the complexity parameter (often denoted by C) reflects the number and seriousness of violations to the margin and hyperplane that we are willing to tolerate.

By using a soft margin, SVM generally perform well in generalizing to fresh instances. SVM has proved to be a powerful method for classification that has outperformed other methods in a wide variety of application domains e.g. text categorization and face and fingerprint identification, and has been successful in various studies of credit risk evaluation [6]. As a result of the SVM depending primarily on (possibly few) support vectors, the method is quite robust to data noise. A weakness of SVM is that the output of the model (weights) are not particularly interpretable.

5.4. Nearest Neighbor

The K-nearest neighbors method classifies an unseen data instance using the classification of the instance(s) that are closest to the unseen instance, using some specified distance metric [15]. The algorithm determines the distance between the test instance and every training example and selects the set of k training instances to the test example then, it predicts the most common class of these nearest training examples by a majority-voting scheme [16]. In the special case where k=1, we simply refer to the algorithm as the nearest neighbor classifier. For simplicity, we use k=1 in the experiments in this paper.

By far the most common distance metric to use with k-nearest neighbors is Euclidean distance. The Euclidean distance D between two data vectors in n-dimensional Euclidean space $\mathbf{x} = (x_1, x_2, , x_n)$ and $\mathbf{y} = (y_1, y_2, , y_n)$ is given by

$$D(x, y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2} \tag{11}$$

An advantage of the nearest neighbor classifier is in its simplicity, in that the algorithm is relatively easy to understand and also can guide ones intuition [17]. Its theoretical basis is appealing when considering how the algorithm is aligned intuitively to how traditional loan decisions were made via human judgment, where credit officers compared the application characteristics to previous similar applicants in order to classify a loan applicant.

Another advantage is that since all work is done when we want to classify a new instance rather than when the training data is processed [12], classification speed is very fast. On the other hand, since the distance of all combinations between test and training instances must be computed, the training speed of the algorithm is generally slow.

5.5. Decision Trees: J48 and Random Forest

The J48 decision tree is the WEKA implementation of the widely used C4.5 decision tree developed by Quinlan [18]. Decision Trees summarize training data in a tree structure, where each branch represents an association between attribute values and a class label [19]. The tree recursively splits the training data based on tests on the potential of feature values in separating the classes (where for C4.5/J48, potential is measured by information gain). New test instances are simply classified by following the decision rules from the top node of the tree downwards. The main advantage of the decision tree classifier is in its output: the tree structure provides a compact and interpretable structure with decision rules to classify fresh test instances. However, if not pruned adequately, trees can become overly complex and can overfit the training data.

Random Forests [21] are an ensemble method, which combine the votes of multiple decision trees that have been generated, through bagging. Candidate features to split on are selected randomly, from a set of the most important features. The random forest model calculates a response variable y by creating many (usually several hundred) different decision trees and then putting each object x to be modeled down each of the decision trees - the response is then determined by evaluating the responses from all of the trees [22]. Although random forests have proven to usually perform better than single decision trees, Random Forests lose their interpretability when compared with a single decision tree structure.

6. Model Evaluation Metrics

6.1. Confusion Matrix

A common way to evaluate the performance of a classifier is by observing the confusion matrix (table 2). A confusion matrix shows all of the instances in the data set, categorized into four different categories: Goods that were classified correctly as such (TP True Positives), Bads that were classified correctly as such (TN True Negatives), Goods that were incorrectly classified

as Bads (FN False Negatives), and Bads that were incorrectly classified as Goods (FP False Positives).

	Classified Bad	**Classified Good**
Actual Bad	TN	FP
Actual Good	FN	TP

Table 2: Confusion Matrix in Credit Scoring

The true positive rate is the proportion of goods that were correctly classified as such:

$$TPRate = \frac{TP}{TP + FN} \tag{12}$$

And similarly, the true negative rate is the proportion of bads that were classified correctly as such:

$$TNRate = \frac{TN}{TN + FP} = 1 - FPRate \tag{13}$$

In the credit scoring setting, the costs of misclassification are quite different. False negatives only refer to the opportunity cost of lost interest that could have been gained, whereas for false positives, the lender loses some or all of not only the interest, but also the principal that was to be repaid [5]. Therefore, false positives are significantly costlier since these are people who were classified as being good by the model, and granted a loan, but actually turned out to be bad.

Accuracy a common metric to evaluate the performance of classifiers. Accuracy represents the proportion of total instances that were classified correctly:

$$Accuracy = \frac{TP + TN}{TP + FP + TN + FN} \tag{14}$$

When the input data set is balanced, and there are equal costs associated with errors, accuracy is an appropriate measure to maximize when selecting a model. However, with imbalanced data and unequal error costs, the ROC curve or other similar techniques are more appropriate [26]. In the case of highly imbalanced data (as we have), it is trivial for a classification model to obtain very high accuracy by simply classifying every loan applicant as a good [25]. But if all loan applicants are classified as good by the model; no bad applicants will be correctly classified as bad, which would obviously be a poor model for our purposes.

6.2. ROC Curve Analysis

An ROC curve plots the true positive (TP) rate the proportion of good risks that are correctly classified, as a function of the false positive (FP) rate, for the whole possible range of cut-off values [28]. It is often easier to think of the x-axis as $1 - TNRate$ (see equation 14). The ROC curve represents a trade-off between TP Rate, and TN Rate. Resampling or cost sensitive learning can move us to a different point on the ROC curve. More favorable models have ROC curves that are situated towards the top-left hand corner of the (TP,FP) space [12].

Rather than attempting to maximize accuracy, in the presence of imbalanced data it is common to attempt to maximize the area-under-ROC-curve (AUC) measure. The AUC measure can show the general performance of a model to another model, in one number.

7. Data

In this section, we outline some of the data preprocessing steps undertaken, as well as the characteristics of the target data set.

7.1. Feature Extraction

Data preprocessing is an important and often time consuming part of any real-world data mining project.

In the present study, data preprocessing firstly involved feature extraction from the lending company's data warehouse. The loan outcome data, along with applicant features from the online application form, were held in one particular table in the data warehouse. The Credit Sense data was held in other separate tables in the data warehouse. Application form features and bank statement features were extracted from the lending company's data warehouse, using a number of OLAP queries in SQL Server. The two sets of features were then joined using the Pandas library in Python.

Relevant features to extract from the bank statement data were decided on consulation with two senior staff members at the lending company, both of whom have deep knowledge of the credit industry.

Location and Merchant information was also initially extracted from the bank statement transactions, but were not found to be useful for modelling purposes (although location information is useful for address verification purposes).

The feature set is shown in table 3 in section 8.

7.2. Class Values: Defining 'Bad'

In credit scoring, we often refer to good risks or simply 'goods', and bad risks or simply 'bads'. It is important to be clear in our definition of what represents the good class and what represents the bad class for the class values in the target data set.

Upon consultation with a number of staff within the lending company, it was decided to use the following as the definition of bad: *a bad is someone who has either been 90 days or more in arrears or has already defaulted.* Note that the 90 days in arrears aspect of this definition is consistent with the definition of default under the Basel II Accord [20].

7.3. Feature Selection

In machine learning, feature selection methods are generally divided into two groups: filter methods and wrapper methods. Wrapper methods are scheme dependent, and evaluate different possible subsets of features from the original set to come up with the best subset for that particular classifier [12]. Due to the fact that they evaluate many different subsets of features using the learning scheme, wrapper methods are computationally expensive and also, they can sometimes result in over fitting. They were not used in this study because of their computational expense. Filter methods on the other hand, look at the predictive power of attributes using a particular measure, in a manner that is independent of the classifier used. Yang, Pedersen [31] and Forman [32] conducted comparative studies on filter methods in the context of text classification, and found that chi-square ranking and information gain ranking are among the best methods of feature selection for classification purposes ([31], [32] as cited in [30]).

The approach in this study for feature selection was to rank the features based on their chi-square statistic and information gain with respect to the class, and vary the number of features that are included in the model, and investigate the AUC value at different numbers of features. It was found that information gain produced very similar rankings to the chi-square statistic, and therefore ultimately it was decided to use only the chi-square statistic as the means of ranking features. Feature selection is performed within the filteredClassifier meta learner in WEKA, so that we do not introduce bias by performing feature selection before training the model.

7.4. Feature Discretization

Numeric variables that are used in credit scoring are usually discretized into bins. Discretizing continuous variables has a number of advantages that are mentioned by [20], including that doing so offers an easier way to deal with outliers, makes it easy to understand the nature and strength of relationships, and enables one to develop insights into the behaviour of risk predictors.

We discretize numeric features using the Discretize filter within the filteredClassifier meta-learner within WEKA, before applying the attribute selection filter. The binning algorithm used is the popular entropy-based multi-interval discretization method proposed by Fayyad & Irani [29].

7.5. Class Imbalance

Data sets used for credit scoring are usually imbalanced, since the number of bad risks is usually small relative to the number of good risks. Our target data set was highly imbalanced, with the proportion of bad instances only 1.6% of total number of instances (7401 good instances, 121 bad instances). There are a number of approaches to dealing with the class imbalance problem. One is using cost-sensitive learning [25], where penalties are imposed for particular types of misclassification. However, this approach relies on knowing the exact cost of misclassifications, which can be difficult to determine.

Another approach is to reduce the number of instances in the good class in some way (undersampling), or to boost the number of instances in the bad class (oversampling). A commonly used approach to oversampling the minority class is called Synthetic Minority Over-Sampling Technique (SMOTE), proposed by Chawla et al [26], which creates synthetic instances for the minority class by taking each minority class instance and introducing synthetic instances along the line segments joining any/all of the k minority class nearest neighbors. Brown and Mues [24] investigated the performance of a number of credit scoring models on real world data sets with different levels of imbalance, and found that random forests performed well with data sets with large imbalance, but SVM with a linear kernel, QDA and C4.5 all performed poorly with imbalanced data. Marques et al [27] reviewed the effectiveness of a number of oversampling techniques (e.g. SMOTE and newer extensions of SMOTE) compared with undersampling techniques, and found that the oversampling techniques outperform the undersampling techniques in most cases, especially when logistic regression is used.

8. Experimentation and Results

8.1. Experimental Setup

Most learning schemes output probabilities as part of the classification process. For example, by default, logistic regression classifies an applicant as good if the conditional probability of being good, given their attribute vector x, is greater than 0.5, and bad otherwise. That is, If $P(Good|\mathbf{x}) > 0.5$ then classify as Good; else classify as Bad.

It was found in the experimentation that using the default cut-off value of 0.5 generally resulted in poor discrimination in the bad class (nearly all-bad risks would be classified as good risks), except for the Naive Bayes classifiers, which provided quite good performance whether the cut-off probability

The threshold selector selects a mid-point threshold on the probability output of a particular classifier. The midpoint threshold altered from the default of 0.5, and is set so that a given performance measure is optimized. In this study, the threshold selector was run to optimize the F1 measure, which balances classifier precision and recall. The F1 measure is defined as the harmonic mean between precision p and recall r [16]:

$$F1 = \frac{2pr}{p+r} \tag{15}$$

(The harmonic mean tends to be closer to the smaller of the two values used in the calculation).

For the purposes of initial comparison of candidate models, all of the base classifiers are wrapped within the thresholdSelector meta-classifier in WEKA for the experiments. Once we have found the best performing model, we will then investigate the use of cost sensitive learning (rather than resampling) to deal with the class imbalance.

The model configurations in WEKA were set up so that base classifier (Nave Bayes, Logistic Regression, SVM etc.) were wrapped within the filteredClassifier meta-learner (which performs attribute discretization and attribute selection), and the filteredClassifier was in turn wrapped within the thresholdSelector meta-learner.

Experimentation was performed with three different feature sets. The two baseline feature sets: the first consisting of the application form features currently used for scoring by the lending company, and the second consisting of the bank statement-derived features. The final feature set is the combined

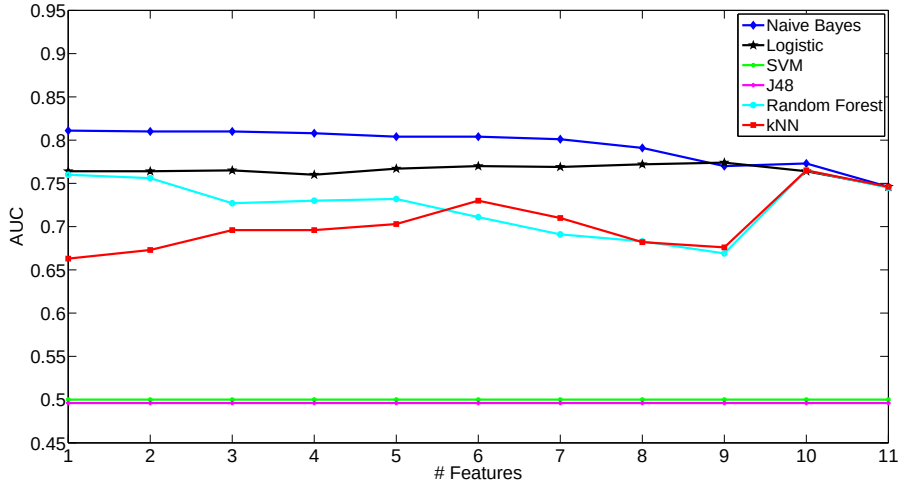

Figure 2: Comparison of AUC for each candidate model with the existing application form features

feature set consisting of both the application form, and the bank statement-derived features.

The features in each feature set were ranked according to their Chi-Square Statistic with respect to the class. The combined feature set rankings are shown below in table 3. All features with a positive Chi-Square Statistic value were initially included in the model, and the lowest ranked features were eliminated from the model one by one. For each feature set, each classifier was run, and the AUC value was recorded.

8.2. Baseline Model with Application Form Features

The AUC for each learning scheme for the specific number of application form features is shown in figure 1. Naive Bayes was found to be the best performing classifier, followed by Logistic Regression. Random Forest and Nearest neighbors yielded similar performance to each other, but were both inferior to Naive Bayes and Logistic Regression. It is noticeable that both the SVM and J48 models perform very poorly with this highly imbalanced data set (consistent with the findings of [24]), with both classifiers being unable to identify any bad instances.

Rank	Feature	Source	Chi-Square
1	Residential Status	Application Form	140.5529
2	Age	Application Form	106.753
3	Marital Status	Application Form	85.0343
4	Loan Purpose	Application Form	82.0138
5	ATM Withdrawal Amount	Bank Statement	79.9832
6	Employment Type	Application Form	55.6844
7	Other Loan Repayments Per Month	Bank Statement	41.9988
8	Employment Duration Years	Application Form	37.638
9	Months at Current Address	Application Form	36.2622
10	Declared Number of Credit Cards Held	Application Form	32.9058
11	Number of Bank Accounts Held	Bank Statement	30.476
12	NZ Region	Application Form	25.2823
13	Visa	Bank Statement	16.987
14	Kiwibank	Bank Statement	12.5209
15	Individual Or Joint	Application Form	11.7183
16	Gender	Application Form	9.5036
17	Mastercard	Bank Statement	5.6001
18	BNZ	Bank Statement	4.722
19	Amex	Bank Statement	4.1162
20	Diners	Bank Statement	1.5046
21	Tax	Bank Statement	0.8387
22	Westpac	Bank Statement	0.7701
23	ANZ	Bank Statement	0.3769
24	Unemployment Benefit	Bank Statement	0.3768
25	TSB	Bank Statement	0.2995
26	Student Allowance	Bank Statement	0.1636
27	ASB	Bank Statement	0.0187
28	Family Allowance Australia	Bank Statement	0.0163
29	Pension	Bank Statement	0.0163

Table 3: Ranking of Combined Feature Set by Chi-Square Statistic with respect to the class. Note that the number of Credit Cards could also be obtained from the bank statement. There were slight differences in the values between the two sources, and the application form variable had slightly higher Chi-Square value compared to the bank statement variable (32.9 vs 22) and was therefore ultimately included in the final feature set

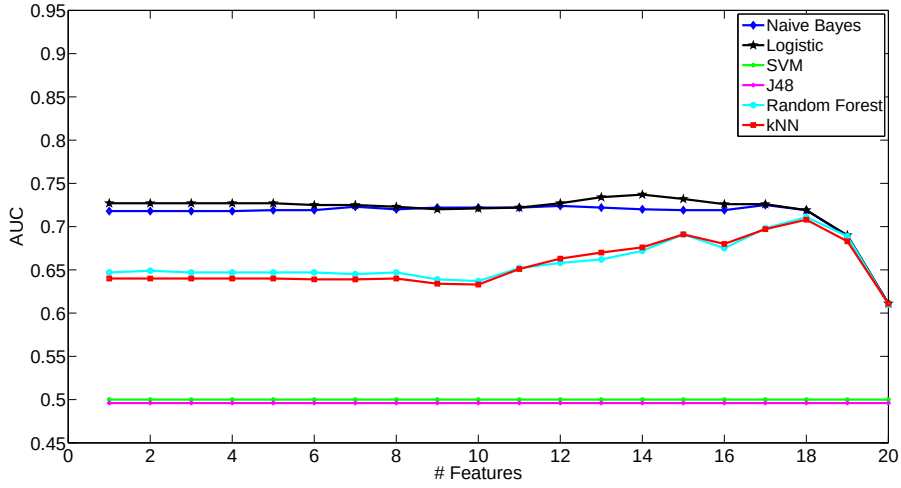

Figure 3: Comparison of AUC for each candidate model with the bank statement-derived features

8.3. Baseline Model with Bank Statement Derived Features

The AUC for each learning scheme for the specific number of bank statement derived features is shown in figure 2. Naive Bayes and Logistic Regression are the best performing classifiers with the bank statement-derived features. Random Forest and nearest neighbor classifier again both had similar performance to each other. Comparing figure 1 and figure 2, we can see that AUC values for the candidate models are clearly higher with the application form feature set than with the bank statement derived feature set. Thus, we can conclude that the lending company should not develop a credit scoring model solely with the bank statement-derived features.

8.4. Combined Feature Set Model

We have established that models performed better with the existing feature set than with the bank statement-derived feature set. We now investigate the extent to which the bank statement-derived features can complement the existing features to improve AUC, and thus the scoring model. The results using the combined feature set are shown in figure 3. We can see that Naive Bayes is clearly the best performing classifier on the combined feature set. The Naive Bayes classifier reaches its highest AUC value (0.829) when 16 features are included in the model. This feature set consists of all of the

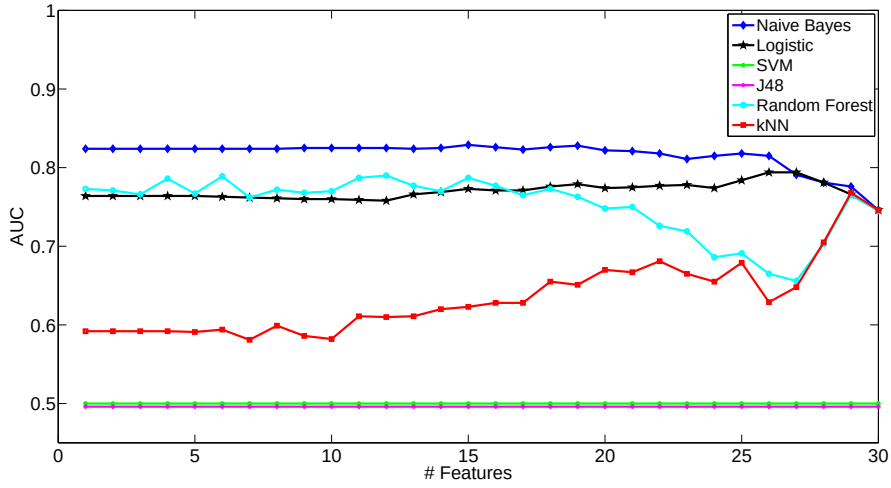

Figure 4: Comparison of AUC for each candidate model with the combined feature set

11 application form features, along with 5 of the new bank statement derived features.

9. Model Evaluation and Interpretation

9.1. Evaluating the Model

We now investigate specifically whether the combined 16-feature Naive Bayes model is superior to a Naive Bayes model based on the existing 11 application form features that are currently used for scoring by the lending company. We can do this by visually comparing the ROC curves for each model in the same trade-off space.

From figure 4 (plotted using the Knowledge Flow interface in WEKA), we can observe that the ROC curve for the combined 16-feature model is generally superior for nearly all trade-off points, except for the regions circled. Thus, we can conclude that the combined feature model provides superior performance to the model based solely on the 11 existing features obtained from the application form, and that the new features from the bank statements have improved the model.

A cost matrix contains the cost of misclassifying a good as a bad (false negative), and the cost of misclassifying a bad as a good. Estimation of these specific costs was not undertaken in the present study, but in any case, we

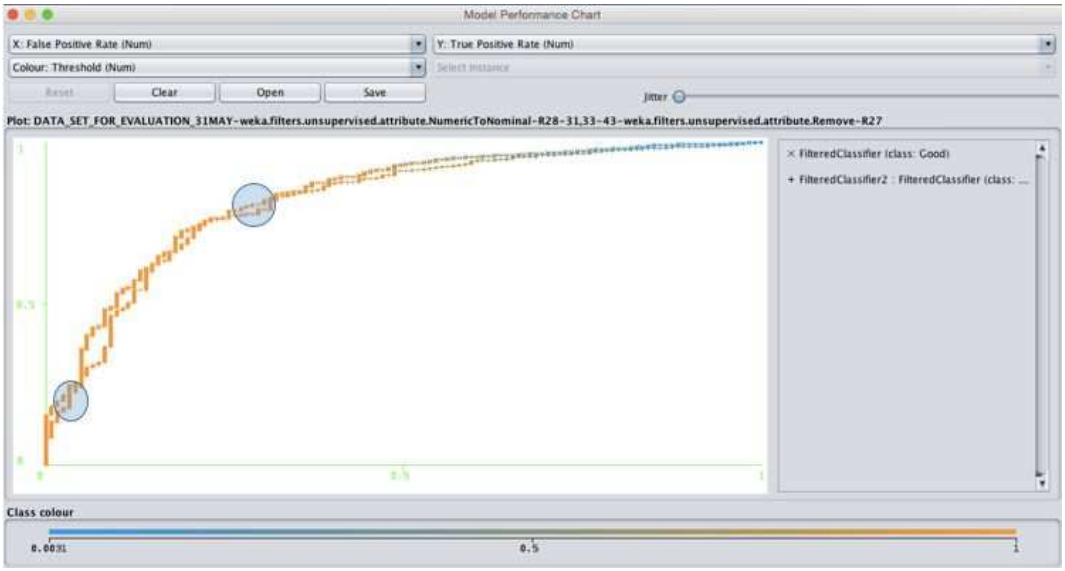

Figure 5: Naive Bayes for the combined feature model (top - FilteredClassifier) vs. Naive Bayes for the traditional feature model (FilteredClassifier2).

can suppose that the cost of misclassifying a good as a bad is $1, and the cost of misclassifying a bad as a good $x; in which case, the cost of misclassifying a bad as a good is x times as great as misclassifying a good as a bad.

Table 4 shows the accuracy, TP rate, the number of goods out of the 7401 total that were classified correctly, the TN rate, and the number of bads out of the 121 total that were classified correctly by the combined 16-feature Naive Bayes model, for six different cost ratios. Note that as we increase the value of x, we can gain in true negative rate and move along the ROC curve, but at the expense of TP rate and overall accuracy.

Further work in conjunction with the lending company is needed to determine the exact costs of misclassification, and therefore to end up on the appropriate trade-off point on the ROC curve.

9.2. Bank Statement Derived Features

As has been mentioned, in the final model, there were 5 features that were derived from the bank statement data, which were found to be of use in complementing the existing application form features. Table 5 shows each of these features, bins and the weight of evidence value for each bin. Note that the weight of evidence represents the natural logarithm of the odds ratio [20]

	1	2	5	10	20	50
Accuracy	0.944	0.928	0.899	0.872	0.839	0.791
True Positive Rate	0.953	0.936	0.904	0.876	0.842	0.793
Goods Classified Correctly (out of 7401)	7059	6931	6700	6493	6239	5876
True Negative Rate	0.413	0.479	0.554	0.603	0.661	0.686
Bads Classified Correctly (out of 121)	50	58	67	73	80	83

Table 4: 16-feature Naive Bayes Model Performance Metrics with different misclassification cost ratios. For example, the second column assumes that a false positive is twice as costly as a false negative.

Large ATM withdrawals may be indicative of impropriety or lack of discipline in spending.

Somewhat surprisingly, higher other loan repayments are associated with a lower risk of default. This is perhaps because it is indicative of higher earning capacity in terms of being able to pay off other loans, particularly a mortgage. Further work in distinguishing between other small loans and other loans such as mortgages could shed more light on this.

People holding more bank accounts are shown to be better risks. Possible reasons for this could be that it is indicative of better handling of their finances, or alternatively indicative of shared finances e.g. joint accounts.

We have also found that holding a Visa card is related to being a good risk, as well as not banking with Kiwibank, which is perhaps a reflection of customer quality.

10. Conclusions and Future Work

This research has supplemented features declared by loan applicants on a loan application form, that were currently used for credit scoring by a New Zealand lending company, with undeclared features that were extracted from bank statements provided by loan applicants. This work has highlighted the potential of using bank statement information through applications such as Credit Sense, to obtain additional features to improve credit scoring.

Although the bank statement features were not sufficiently predictive by themselves, our experimental results showed that a number of the bank statement derived features have value in supplementing the application form features, in improving a credit scoring model.

Attribute	**Attribute Value**	**WOE**
Bank with Kiwibank	No	0.159
	Yes	-0.568
Hold a Visa Card	No	-0.267
	Yes	0.585
Number of Bank Accounts Held	0 to 2	-0.399
	3 or more	0.703
Other Loan Repayments per Month	Under $445 (low)	-0.514
	$445 or more (high)	0.719
Amount withdrawn from ATM Machines	Under $2131 (low)	0.412
	$2131 or more (high)	-1.118

Table 5: Weight of evidence values for each of the 5 bank statement-derived attribute bins, in the final Naive Bayes combined feature model. Numeric features were discretized using the Fayyad and Irani method [29]

.

In particular, we found that a Naive Bayes model that used 16 combined features provided the best performance in terms of area under the ROC curve. This 16 feature model consisted of all 11 application form features that are currently used for credit scoring by the lending company, along with 5 new features derived from the bank statements. These 5 features were: amount withdrawn from ATM machines, other loan repayments per month, whether they bank with Kiwibank, whether they hold a Visa card, and the number of bank accounts held.

There are a number of avenues for further work. One area is to extend this work is by incorporating reject inference methods. This involves attempting to assign a value to rejected applicants or determining the status of each reject and then using this value in the target data set [33]. This was not possible in the present study due to the non-availability of this type of data.

This work has also highlighted the potential of Naive Bayes to be used with extremely imbalanced data sets. It is interesting to note that Naive Bayes was not considered to be a commonly applied model in the credit scoring domain by [6]. We intend to perform further investigation of the ability of Naive Bayes to be used in the credit scoring context, and specifically with highly imbalanced data sets.

Another extension would be to further investigate the costs of misclassification with the lending company. This would involve coming up with

accurate estimates of the cost of false negatives, and the cost of false positives. In doing so, the lending company would be moving to the appropriate trade-off point on the ROC curve.

From an organizational perspective, further work would be required in implementing the credit scoring model in practice e.g. integrating the model with the data warehouse and other IT systems, testing and so on.

In a separate line, data from bank statements obtained through applications like Credit Sense could be used in investigating possible fraud cases, for example, by using methods of outlier detection.

References

References

[1] Thomas LC, Edelman DB, Crook JN. Credit scoring and its applications. Siam; 2002.

[2] Rosenberg E, Gleit A. Quantitative methods in credit management: a survey. Operations research. 1994 Aug; 42(4):589-613.

[3] Hand DJ, Henley WE. Statistical classification methods in consumer credit scoring: a review. Journal of the Royal Statistical Society: Series A (Statistics in Society). 1997 Sep 1; 160(3):523-41.

[4] Thomas LC. A survey of credit and behavioural scoring: forecasting financial risk of lending to consumers. International journal of forecasting. 2000 Jun 30;16(2):149-72.

[5] Abdou HA, Pointon J. Credit scoring, statistical techniques and evaluation criteria: A review of the literature. Intelligent Systems in Accounting, Finance and Management. 2011 Apr 1; 18(2-3):59-88.

[6] Yu L, Wang S, Lai KK, Zhou L. BioInspired Credit Risk Analysis. Springer; 2008.

[7] Fayyad U, Piatetsky-Shapiro G, Smyth P. From data mining to knowledge discovery in databases. AI magazine. 1996 Mar 15; 17(3):37.

[8] Pyle D. Data preparation for data mining. Morgan Kaufmann; 1999.

[9] Krzysztof JC, Pedrycz W, Swiniarski RW, Kurgan LA. Data Mining: A Knowledge Discovery Approach. Springer. 2007.

[10] Thomas LC. Consumer Credit Models: Pricing, Profit and Portfolios: Pricing, Profit and Portfolios. OUP Oxford; 2009.

[11] James G, Witten D, Hastie T, Tibshirani R. An introduction to statistical learning. New York: springer; 2013 Feb 11.

[12] Hall M, Witten I, Frank E. Data mining: Practical machine learning tools and techniques. Kaufmann, Burlington. 2011.

[13] Cortes C, Vapnik V. Support-vector networks. Machine learning. 1995 Sep 1; 20(3):273-97.

[14] James G, Witten D, Hastie T, Tibshirani R. An introduction to statistical learning. New York: springer; 2013 Feb 11.

[15] Bramer M. Principles of data mining. London: Springer; 2007 Mar 6.

[16] Pang-Ning T, Steinbach M, Kumar V. Introduction to data mining. In Library of congress 2006 (Vol. 74).

[17] Aha DW, Kibler D, Albert MK. Instance-based learning algorithms. Machine learning. 1991 Jan 1; 6(1):37-66.

[18] Quinlan JR. C4. 5: Programming for machine learning. Morgan Kauffmann. 1993 Jan:38.

[19] Novakovi J, trbac P, Bulatovi D. Toward optimal feature selection using ranking methods and classification algorithms. Yugoslav Journal of Operations Research ISSN: 0354-0243 EISSN: 2334-6043. 2011 Jan 1; 21(1).

[20] Siddiqi N. Credit risk scorecards. Developing and implementing intelligent credit scoring. SAS Institute Inc. New Jersey. USA. 2006.

[21] Breiman L. Random forests. Machine learning. 2001 Oct 1; 45(1):5-32.

[22] Horning N. Random Forests: An algorithm for image classification and generation of continuous fields data sets. In Proceedings of the International Conference on Geoinformatics for Spatial Infrastructure Development in Earth and Allied Sciences, Osaka, Japan 2010 Dec 9 (Vol. 911).

[23] Chapman P, Clinton J, Kerber R, Khabaza T, Reinartz T, Shearer C, Wirth R. CRISP-DM 1.0 Step-by-step data mining guide.

[24] Brown I, Mues C. An experimental comparison of classification algorithms for imbalanced credit scoring data sets. Expert Systems with Applications. 2012 Feb 15;39(3):3446-53.

[25] Elkan C. The foundations of cost-sensitive learning. InInternational joint conference on artificial intelligence 2001 Aug 4 (Vol. 17, No. 1, pp. 973-978). LAWRENCE ERLBAUM ASSOCIATES LTD.

[26] Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: synthetic minority over-sampling technique. Journal of artificial intelligence research. 2002;16:321-57.

[27] Marqus AI, Garca V, Snchez JS. On the suitability of resampling techniques for the class imbalance problem in credit scoring. Journal of the Operational Research Society. 2013 Jul 1;64(7):1060-70.

[28] Bastos J. Credit scoring with boosted decision trees. 2007.

[29] Fayyad UM, Irani KB. On the handling of continuous-valued attributes in decision tree generation. Machine learning. 1992 Jan 1;8(1):87-102.

[30] Geng X, Liu TY, Qin T, Li H. Feature selection for ranking. InProceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval 2007 Jul 23 (pp. 407-414). ACM.

[31] Yang Y, Pedersen JO. A comparative study on feature selection in text categorization. InICML 1997 Jul 8 (Vol. 97, pp. 412-420).

[32] Forman G. An extensive empirical study of feature selection metrics for text classification. Journal of machine learning research. 2003;3(Mar):1289-305.

[33] Thomas, L. C. (2009). Consumer Credit Models: Pricing, Profit and Portfolios: Pricing, Profit and Portfolios. OUP Oxford.